

EXHIBIT G



U.S. Copyright Office

Notice of Inquiry on Artificial Intelligence & Copyright (Dkt. 2023–6)

Comments of Meta Platforms, Inc.

October 30, 2023

I. Introduction¹

Imagine a world where language barriers evaporate and people can communicate seamlessly in real time; where governments, the private sector, and civil society can come together to find solutions to the challenges of climate change before it's too late; where vaccines can be developed and deployed before a pandemic takes grip of humanity. All of this—and so much more—is possible with the power of AI. And all of these opportunities could be squandered if the delicate balance between copyright, and innovation and competition is miscalibrated.

Meta Platforms, Inc. builds technology that helps people connect and share, find communities, and grow businesses. As a creative rightsholder, one of the world's leading technology innovators, and a pioneer in the field of artificial intelligence, Meta is uniquely positioned to provide its views on the issues raised in the Office's Notice of Inquiry on Artificial Intelligence and Copyright. 88 Fed. Reg. 59942. Meta appreciates the Office's attention to these important issues and the opportunity to submit these comments for the Office's consideration.

Meta is committed to making sure that the vast opportunities unlocked by this new technology are open to everyone. To make sure that happens, Meta has taken on a unique role in the AI landscape. Meta continues to release highly valuable models and tools to the public, free of charge. This includes tools for machine translation,² computer vision,³ fairness evaluation,⁴ and, most recently, several versions of a large language model called Llama.⁵ The reason for doing this is simple: by openly sharing its research and breakthroughs with the entire AI community, Meta can reduce the financial and technical barriers to AI development and empower a diverse field of developers, researchers, and businesses to advance the field of AI as a whole. Meta is poised to be a leader in bringing these tools to people around the world.

Meta is also taking a leadership role in ensuring the responsible development of AI technologies, guided by five pillars of responsible AI: (i) protecting the privacy and security of people's data; (ii) ensuring that everyone is treated fairly when using our products, and that our

¹ This section addresses questions 1 and 2.

² See Meta, "200 Languages Within A Single AI Model: A Breakthrough in High-Quality Machine Translation," Meta Blog, available at: <https://ai.meta.com/blog/nllb-200-high-quality-machine-translation/>.

³ See Meta, "Introducing Segment Anything: Working Towards the First Foundation Model for Image Segmentation," Meta Blog (Apr. 5, 2023), available at: <https://ai.meta.com/blog/segment-anything-foundation-model-image-segmentation/>.

⁴ See Meta, "Introducing Casual Conversations v2: A More Inclusive Dataset to Measure Fairness," Meta Blog (Mar. 9, 2023), available at: <https://ai.meta.com/blog/casual-conversations-v2-dataset-measure-fairness/>.

⁵ See Meta, "Meta and Microsoft Introduce the Next Generation of Llama," Meta Newsroom (July 18, 2023), available at: <https://about.fb.com/news/2023/07/llama-2/>.



products work equally well for everyone; (iii) developing AI systems that meet high performance standards, including by testing them to ensure they behave safely and as intended; (iv) providing our users with more transparency and control around how data about them is collected and used; and (v) building reliable processes to ensure accountability for our AI systems and the decisions they make.

To assist the Office in its AI initiative, this document opens with an overview of the general technical process underlying Generative AI models,⁶ with a particular focus on large language models and how they are being deployed in today’s AI landscape. *See infra* Part II. It then highlights Meta’s unique role in this ecosystem. *See infra* Part III. Next, it discusses how the use of copyrighted material to train Generative AI models is a non-consumptive use that does not trigger the rights protected by copyright, and even if it did, it would be a fair use. Indeed, there is no legitimate market through which AI developers might pay rightsholders to extract unprotectable facts, ideas, and concepts from their works, which means this use is quintessentially fair. *See infra* Part IV. Finally, it identifies several problems with the proposals to create statutory licensing mechanisms, *see infra* Part V, and briefly addresses the longstanding and salutary principle that copyright law does not and should not protect artistic styles, *see infra* Part VI.

II. AI Models and Their Function

A. Training Models is a Statistical Process for Learning Facts About the World⁷

Today’s Generative AI models are undoubtedly impressive, but the broad principles on which they rely are somewhat basic. The goal of training a Generative AI model is fundamentally the same as the goal of building any statistical model: to derive patterns and relationships from a set of existing data. A meteorologist who wants to build a model to predict the weather, for example, needs a dataset of past weather data from which the model can derive statistical information representing the complex relationship between air pressure, humidity, temperature, and other factors, and their impact on the likelihood of (for example) a thunderstorm. In the same way, an AI developer who wants to build a language model needs a dataset of text from which the model can derive statistical information reflecting the relationships between different words and the grammatical rules which govern speech patterns and language.

For both the meteorologist and the AI developer, the *quantity* and *diversity* of the data is critical. A meteorological model based on only a few days of weather data will be less useful and effective than a model based on 50 years of weather data. And a model based on weather data from the month of June will not accurately predict the likelihood of rain in the middle of December. Similarly, an AI language model based on a small amount of data will not have enough information from which to derive the highly complex and nuanced rules, meanings, and contradictions of human language. And a model trained on a narrow category of works—like scientific journals—will not have enough information to accurately discern the patterns of modern, everyday human speech. To build a model that can realistically emulate all facets of human language, the developer

⁶ “Generative AI” refers to applications of AI technology to generate outputs in the form of expressive material such as text, images, audio, or video. *See* 88 Fed. Reg. 59948.

⁷ This section generally addresses question 7.



needs a collection of data that includes a very large number of examples, reflecting a broad range of speech—from casual banter, to literary prose, to scientific jargon.

In both circumstances, it is the analysis of a set of data *as a whole*—rather than any individual piece of data—that provides value. Whether a meteorological model includes or excludes a single row of data regarding the barometric pressure, temperature, and humidity on a particular day at a particular location will not materially affect the model’s effectiveness. This is because the model’s function is to identify the statistical *patterns* that play out across a broad range of different data points. Similarly, whether a particular block of text is included or excluded from a large language model’s training data makes little difference. The point of model training is not to simply reproduce pieces of text, but rather to observe and identify the statistical *patterns* across many, many examples of text.


B. How Large Language Models “Learn”⁸

The primary difference between a traditional statistical model and a more complex large language model is the way in which the latter stores and organizes information. Large language models rely on a variety of sophisticated mathematical techniques and algorithms, many of which have been developed in the past few years.⁹ But each of these techniques are, at their core, methods for deriving broad patterns from training data and recording those patterns—at varying levels of abstraction—in numbers that computers can process through mathematical operations. In other words, the goal of these techniques is not to extract particular expressive content, but to extract syntactical, structural, and linguistic information from a corpus of works as a whole. We describe an example of such a technique below.

One of the critical ingredients to developing a large language model is “word embedding” or “vectorization,” which is a process of capturing and recording the semantic meaning of individual words. This process begins by breaking blocks of text into fragments or “tokens,” which are either words or parts of words. The goal of word embedding is to create what is essentially a “dictionary” of words that a machine can easily use and understand. AI scientists recognized that one can create such a dictionary by defining words in relation to each other using numerical notations. To take a simple example, imagine you are trying to define the word “aunt.” English speakers understand, through our own education, that “aunt” has a semantic meaning close to “woman” and “sister,” farther from “man” and “brother,” and even farther from “dog,” and still even farther from “chair.” While we understand those points qualitatively, that basic insight can also be expressed mathematically, by plotting the words against each other in a multidimensional space. What AI scientists discovered is that, given enough training data, you can train an AI model to accurately place all the words in the English language into such a “semantic” space, so that the machine can understand, through numbers, what “aunt” means. To provide a simplified example:

⁸ This section generally addresses question 7.

⁹ See, e.g., Ashish Vaswani et al., “Attention Is All You Need” (Aug. 2, 2023), available at: <https://arxiv.org/pdf/1706.03762.pdf>.



In the above example, the words are mapped in a three-dimensional space, and their place within that space is defined with a set of numbers—called a “vector”—with each number reflecting the distance along one of the three axes that define that space, and each axis reflecting some innate characteristic of the word.¹⁰ In reality, however, developing a sufficiently fine-grained semantic understanding of words requires many more dimensions. A “50-dimensional” model, for example, will have 50 numbers in each word’s vector. Below is an example of what a token (“aunt”) and its corresponding vector might look like in an 8-dimensional model:

token		vector						
aunt	2.6	2.6	2.9	1.5	2.5	2.0	0.8	1.2

Before the training process begins, each number in each vector will be assigned a random value, as shown in the example below:

¹⁰ Conceptually, one can think of the axes as reflecting characteristics like how many legs the subject has, how large it is, etc. But in practice, it is the learning algorithm itself that organically determines how to use each vector dimension to capture semantic meaning, and algorithms generally do so in a way that is not easily interpretable by humans.



aunt	2.6	2.6	2.9	1.5	2.5	2.0	0.8	1.2
machine	1.7	2.5	3.0	0.6	1.7	1.5	0.8	1.3
factory	0.7	2.6	2.5	0.9	1.2	0.5	0.4	2.7
plant	1.1	0.7	0.6	1.7	1.0	1.0	2.5	0.9
garden	2.6	2.9	1.9	1.4	2.7	0.7	2.7	1.5
flower	1.6	2.9	1.4	3.0	0.6	1.5	0.7	1.5
susan	1.6	0.7	1.7	1.6	1.0	2.0	0.2	0.8

The process of “embedding” consists of using training data to gradually adjust these random numbers so that they eventually capture the words’ meanings in a numerical way. Suppose, for example, that a model encounters the following sentence in its training data:

Susan’s aunt planted the flowers in the garden.

In processing that sentence, the model would adjust the numbers in the vectors for the words “Susan,” “aunt,” “plant,” “flower,” and “garden” to reflect the fact that those words appeared together in a sentence. After encountering thousands of other sentences or sentence fragments containing some combination of the words “plant,” “flower,” and “garden”—and far fewer sentences combining those words with “Susan” or “aunt”—the model’s vectors will begin to reflect the semantic proximity (or distance) of those words.

Similarly, after the model encounters many sentences or sentence fragments using the words “machine,” “factory,” and “plant” together or in similar contexts, the vector numbers will also reflect this distinct meaning of the word “plant” by reflecting its semantic proximity to this different set of words. Below is an example of what the resulting vectors might look like:

aunt	2.6	2.6	2.9	1.5	2.5	2.0	0.8	1.2
machine	1.3	1.0	0.2	1.4	4.3	2.3	3.3	2.3
factory	1.2	2.1	0.1	3.4	1.0	2.4	3.4	4.4
plant	1.1	3.4	0.1	0.3	1.7	1.6	3.5	5.5
garden	0.2	3.5	4.0	0.3	1.8	3.4	1.2	0.2
flower	0.5	3.3	3.4	0.2	2.3	4.2	2.0	1.3
susan	1.6	0.7	1.7	1.6	1.0	2.0	0.2	0.8

As illustrated here, the semantic proximity of the words “machine,” “factory,” and “plant” is represented by the proximity of the numbers in the first, third, and seventh vector spaces (highlighted in blue). And the semantic proximity of the words “plant,” “garden,” and “flower” is represented by the numbers in the second and fourth vector spaces (highlighted in green). In this



way, the model’s vector for the word “plant” simultaneously reflects both possible meanings of that word. After examining billions of similar sentences—and making billions upon billions of tiny adjustments to the numbers in these vectors—the word vectors eventually become a semantic map of the language, accurately reflecting the meaning of words by encoding the semantic relationship between them.

This simple example demonstrates why both *quantity* and *diversity* of training data inform the quality of a model’s output. A model that has access to only a limited amount of data—*i.e.* the single sentence “*Susan’s aunt planted the flower in the garden*”—knows only that those eight words can occur together in a sentence. But it has no way of knowing that the words “planted,” “flower,” and “garden” have a closer semantic relationship to each other than to the other words in the sentence, like “aunt,” or “susan.” Only by seeing enough other sentences that include different combinations of these words can the model derive enough information to “know” the words “plant” and “garden” are more closely related than the words “garden” and “susan.”

Similarly, large language models require a very *diverse* set of training data to discern the many possible meanings of a single word. A set of training data consisting entirely of manuals, industrial specifications, and patent applications may discern the semantic proximity of the words “plant” with the words “machine” and “factory,” but might be entirely ignorant of the relationship of the words “plant” to the words “flower,” and “garden.” Again, it is the quantity and diversity of the dataset—not the inclusion or exclusion of any individual sentence or any individual source text, *i.e.* “*Susan planted the flower in the garden*”—that allows the model to develop an accurate understanding of the meaning of each word.

This example also illustrates an important fact that is widely misunderstood: an AI model does not “store” or contain copies of the data itself. Rather, the model uses its training data to guide the process of gradually adjusting its representations of abstract meaning—*i.e.*, the vectors in the example shown above. To be sure, the inclusion of the sentence “*Susan planted the flower in the garden*” may have had some miniscule effect on the final value of the numbers in the vectors for those seven words. But the resulting vector values are not determined by any one piece of training data; instead, they reflect the influence of *billions* of other sentences including the same words. That has two critical implications:

- a) no individual piece of content has a particular influence over the vectors of a model;
- b) while a piece of data may assist a model in better understanding the meaning and inter-relationship of words, the model does not copy the particular training data, but rather, reflects the vectors it has learned based on the training set as a whole.

The process of creating a large language model is much more complex than the process described above. Language models capture far more than the meaning of individual words—they also capture grammatical structures, idioms and turns of phrase, facts about the world, etc. They do so through highly complex networks of nodes—comprising what’s known as a “neural network”—which independently identify patterns at various levels of abstraction.

While foundation models contain a vast amount of statistical information regarding language patterns, they are not able to answer user queries or engage in conversations. A pre-



trained model, for example, can be used to predict that the sequence “red, orange, yellow, green, blue” is most likely to be followed by “indigo, violet.” But it cannot answer the question: “what are the colors of the rainbow?” To build a conversational AI agent (*i.e.*, to interact with users and respond to their requests), a pre-trained model must be “fine-tuned” through a variety of techniques. This fine-tuning often involves the use of supervised fine-tuning data (including, for example, curated sets of questions and responsive answers) and “Reinforcement Learning with Human Feedback” or “RLHF,” through which developers use human-created annotations on model responses to further align the model’s behavior to human preferences.¹¹

Importantly, the way a language model “learns” from these various steps of development is similar to the “word embeddings” process described above. While a model might be trained with or without human supervision, using a variety of different uncategorized or categorized data inputs, the result and purpose of all forms of model training is to internalize in the model patterns that reflect language, speech, knowledge, and behavior.

Further, these “fine-tuned” models are most often supplemented with user interfaces (and interface-level controls such as prompt or output filters) that further adjust the model’s behavior. Every user-facing chatbot available today—including, e.g., Meta’s AIs, ChatGPT, Google Bard, and Microsoft’s Bing Chat—features a carefully-designed user interface which includes layers of quality control and safety checks to limit what kinds of queries the model responds to and adjust its outputs accordingly.

Moreover, while the above discussion relates specifically to large language models, the basic point is also true of other kinds of Generative AI models, whether video, image, or multimodal models. The goal of such models is to simply extract non-expressive facts and statistics from training data (e.g., what characteristics typify a cat) and use them to generate new content (e.g., an entirely new picture of a cat).

C. Generative AI is a Tool That Enhances Rather Than Supplants Human Creativity and Productivity, and is Ultimately One That People Control¹²

The role of a Generative AI model is to enhance human creativity and productivity. Generative AI is a tool, no different from the printing press, the camera, or the computer. Those technologies changed the nature of creative endeavors, and were a tremendous boon to human creativity and productivity. Generative AI will be no different.

Generative AI tools are still in their infancy, but we are already seeing many examples of users and creators using Generative AI tools in creative, interesting, and useful ways.

- Latimer is a LLM founded and led by entrepreneur John Pasmore that builds on Llama 2 and other AI models by adding books, oral histories, and other material from

¹¹ See generally Meta, “Llama 2: Open Foundation and Fine-Tuned Chat Models” §§ 3.1, 3.2 (July 19, 2023), available at: <https://arxiv.org/pdf/2307.09288.pdf>.

¹² This section generally addresses questions 1 and 25.



historically underrepresented communities with the goal of making a more inclusive Generative AI tool.¹³

- Architectural designer Tim Fu¹⁴ is integrating Generative AI into visionary design; as one example, he used AI to design the Chromasis armchair sold by Italian luxury furniture brand Mavimatt.¹⁵
- HiiiWAV¹⁶ is a creator incubator focused on Black artist-entrepreneurs that teaches artists how to build and use AI tools in their craft.¹⁷
- Straightlabs specializes in the development of digital educational platforms and virtual interactive learning environments. Leveraging Meta's AI technology, they empower 3D avatars to offer personalized mentoring and precise responses to inquiries on specific subjects, thereby elevating the overall learning experience and meeting the unique needs of learners.¹⁸
- In the e-commerce space, Shopify used Llama 2 to help build Sidekick,¹⁹ an AI assistant built right into Shopify, that understands all things Shopify and specific context about the merchant. The assistant is able to step in and support merchants with tasks that they are working on, answer questions, fill out forms, edit themes, and even provide analytics.
- In less than a year since they were released, Llama and Llama 2 have been the basis of significant adoption and use by platforms, innovators, and developers. Major platforms such as AWS, Google Cloud, and Microsoft Azure have embraced Llama models on their platforms, while thousands of innovators and startups including Anyscale, Replicate, Snowflake, LangSmith, Scale AI, and so many others are making Llama the foundation for their Generative AI product innovation. The open source and developer community are building a wide variety of tools based on Llama, while major hardware platforms like AMD, Intel, NVIDIA, and Google have boosted the performance of Llama 2 through hardware and software optimizations.²⁰

¹³ See "The Black GPT: Introducing the AI Model Trained with Diversity and Inclusivity in Mind," People of Color in Tech (Oct. 20, 2023), available at: <https://peopleofcolorintech.com/articles/the-black-gpt-introducing-the-ai-model-trained-with-diversity-and-inclusivity-in-mind/>.

¹⁴ See Studio Tim Fu, "Studio Tim Fu: Generative AI in Architectural Design," STF Home Page, available at: <https://www.timfu.com>.

¹⁵ See "Chromasis the World's First Armchair Created with AI," Mavimatt, available at: <https://mavimatt.com/product/chromasis-the-world-s-first-armchair-created-with-ai/>.

¹⁶ See HiiiWAV, "HiiiWAV Presents AFRO AI," HiiiWAV Home Page, available at: <https://hiiwav.org>.

¹⁷ See HiiiWAV, "Afro AI Workshop Sign Up," Eventbrite, available at: <https://www.eventbrite.com/e/afro-ai-workshop-sign-up-tickets-711449413347>.

¹⁸ See Straightlabs, "Education Meets Technology," Straightlabs Home Page, available at: <https://www.straightlabs.com/>.

¹⁹ See Shopify, "AI Designed for Commerce," Shopify Home Page, available at: <https://www.shopify.com/magic>.

²⁰ See Meta, "The Llama Ecosystem: Past, Present, and Future," Meta Blog (Sept. 27, 2023), available at: <https://ai.meta.com/blog/llama-2-updates-connect-2023/>.



Moreover, as with any tool, people ultimately are responsible for how to use that tool. Just as a photographer chooses the settings on a camera and the subjects she wishes to photograph, the person using a Generative AI model is the one that provides the prompts to the model that determine the content of the output, and then is the one that decides how to use that output. The singular purpose of these models is to enable people to create new creative output to suit their preferences and needs. In that sense, a Generative AI model is not different from other neutral commercial technologies accepted under copyright law, from sound mixing to digital photo editing tools and beyond.

III. Meta's Role in the AI Landscape²¹

Meta's approach to AI has been to make this revolutionary technology available to as many people as possible, as quickly as possible through open innovation. For that reason, it has released its Llama model to the public: first for research uses, and, now, for commercial use as well. Meta was one of the first companies to provide a large language model to researchers when it released OPT-175B in May 2022 on an open-source basis.²² Meta's release enabled researchers to further explore how these language models work, helping further the progress of AI science.

This year, Meta carried forward this open innovation approach by releasing a series of "Llama" models, free of charge for research use, and, with its more recent "Llama 2" model, commercial use as well. This included both "pre-trained" models—which others can fine-tune as they wish—and "fine-tuned" models developed using state-of-the-art reinforcement learning and other methods.²³ Roughly one month later, Meta released yet another model, called Code Llama, fine-tuned to help developers generate computer code.²⁴

These releases have been valuable contributions to the further development of the AI technology ecosystem. They allow researchers, developers, and businesses to leverage the power of AI without having to invest to "pre-train" their own foundation models. To name just a few examples, Meta's open innovation AI releases have enabled researchers to develop novel methods for watermarking model outputs so that they can be identified as AI generated,²⁵ allowed implementers to fine-tune foundation models for as little as \$300 to create their own AI products,²⁶

²¹ This section addresses questions 1 and 2.

²² See Meta, "Democratizing Access to Large-Scale Language Models with OPT-175B," Meta Blog (May 3, 2022), available at: <https://ai.meta.com/blog/democratizing-access-to-large-scale-language-models-with-opt-175b/>.

²³ See *supra* note 11.

²⁴ See Meta, "Introducing Code Llama, A State-of-the-Art Large Language Model for Coding," Meta Blog (Aug. 24, 2023), available at: <https://ai.meta.com/blog/code-llama-large-language-model-coding/>.

²⁵ See John Kirchenbauer, et al., "A Watermark for Large Language Models," 1, 7 (June 6, 2023), available at: <https://arxiv.org/pdf/2301.10226.pdf> (noting reliance on Meta's "Open Pretrained Transformer (OPT) family" to test "a watermarking framework for [] language models").

²⁶ See Vicuna Team, "Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90% * ChatGPT Quality," The Large Model Systems Organization (Mar. 30, 2023), available at: <https://lmsys.org/blog/2023-03-30-vicuna/> ("We introduce Vicuna-13B, an open-source chatbot trained by fine-tuning LLaMA on user-shared conversations . . . The cost of training Vicuna-13B is around \$300.").



and spurred medical research advances by combining language models with scientific databases.²⁷ Meta’s open innovation approach improves the safety and reliability of AI models, by allowing a broader community of developers to test and improve the models. It also creates greater transparency by allowing developers and researchers to uncover ways in which the model is not performing as designed.

IV. Copyright and AI²⁸

Given the profound implications for the future development and deployment of transformative AI innovations, it is critical for policymakers to think carefully about the legal framework that will govern this technology. Copyright is an essential component of that legal framework calibration. The United States is the global leader in innovation and AI development. This is due, in part, to the careful balance our copyright law strikes between rightsholders’ interests and the freedom to innovate with new technologies. Under this established framework, training and development of AI models does not trigger the rights protected by the Copyright Act; and, even if it did, it would be a fair use for reasons elaborated in section IV.B below.

A. Copyright Does Not Prohibit Non-Consumptive Uses

The process of training and developing AI models does not necessarily trigger the rights that copyright exists to protect. The Act grants to copyright holders several limited “exclusive” rights, enumerated in Section 106. The Act does not, however, “give a copyright holder control over all uses of his copyrighted work,” and one who “puts the work to a use not enumerated [by the Act] does not infringe.”²⁹ Moreover, Section 102(b) of the Act provides that “in no case does copyright protection . . . extend to any idea, procedure, process, system, method of operation, concept, principle, or discovery, regardless of the form in which it is described, explained, illustrated, or embodied in [a] work.”³⁰

At a high level, the use of copyrighted content to train AI models does not implicate the rights protected by the Act. As the Supreme Court explained in *Eldred v. Ashcroft*, “every idea, theory, and fact in a copyrighted work becomes instantly available for public exploitation at the moment of publication.”³¹ The extraction of unprotectable facts and ideas from copyrighted works is not itself an infringement of copyright, whether that extraction is accomplished by a human being (by, for example, learning from a book) or by a technological process like the one described in Section II above. It is well established that the extraction of “statistical information”—like “word frequencies, syntactic patterns, and thematic markers [] to derive information on [] nomenclature, linguistic usage, and literary style”—does not implicate the interests protected by

²⁷ See generally Chaoyi Wu et al., “PMC-LLaMA: Towards Building Open-source Language Models for Medicine” (Aug. 25, 2023), available at: <https://arxiv.org/pdf/2304.14454.pdf>.

²⁸ This section generally addresses question 8.

²⁹ *Twentieth Century Music Corp. v. Aiken*, 422 U.S. 151, 154-55 (1975); see also *CDK Glob. LLC v. Brnovich*, 16 F.4th 1266, 1276 (9th Cir. 2021) (“[T]he Copyright Act does not provide copyright owners the exclusive right to use their works.” (emphasis in original)).

³⁰ 17 U.S.C. § 102(b).

³¹ 537 U.S. 186, 219 (2003).



the Copyright Act.³² To the extent the process of AI training consists entirely of the extraction of such information from training data, that process does not violate copyright protections.

Similarly, AI models themselves cannot be fairly characterized as “derivative works” of the training data such that their creation violates the derivative-work right. As the Second Circuit explained in the *Google Books* case, the “derivative works over which the author of the original enjoys exclusive rights ordinarily are those that re-present the protected aspects of the original work, *i.e.*, its expressive content.”³³ For that reason, because the Google Books service under review in that case did not “allow access in any substantial way to” the “expressive content” in the plaintiffs’ books, the Court held that the service could not itself be a derivative work.³⁴ The same principle applies to AI models, which, when functioning properly, do not in any way “re-present” or “allow access” to expressive content on which the model may have been trained.

B. Fair Use and AI

As the Office knows, one of the most critical elements of the balance between rightsholder interests and innovation is the doctrine of fair use. Courts have applied that doctrine to lay the legal groundwork for revolutionary new technologies like internet search,³⁵ while declining to sanction more exploitative technologies, like unauthorized file sharing and unlicensed media clip services.³⁶

The American AI industry is built in part on the understanding that the Copyright Act does not proscribe the use of copyrighted material to train Generative AI models. That understanding flows directly from the fact that model training is a quintessentially non-exploitive use of training material. As explained above, the purpose and effect of training is not to extract or reproduce the protectable expression in training data, but rather to identify language *patterns* across a broad body of content. Doing so does not implicate any of the legitimate rightsholder interests that copyright law exists to protect.

To that end, model training is squarely protected by the fair use doctrine, whose central purpose is to “avoid rigid application of the copyright statute when . . . it would stifle the very creativity which that law is designed to foster.”³⁷ Fair use is a “limitation[.]” on the “exclusive rights” of the copyright holder.³⁸ In 1976, Congress codified the fair use doctrine in Section 107

³² *Authors Guild v. Google, Inc. (Google Books)*, 804 F.3d 202, 209 (2d Cir. 2015).

³³ *Id.* at 225.

³⁴ *Id.* at 226 (“Nothing in the statutory definition of a derivative work, or of the logic that underlies it, suggests that the author of an original work enjoys an exclusive derivative right to supply information about that work of the sort communicated by Google’s search functions.”).

³⁵ *See, e.g., Kelly v. Arriba Soft Corp.*, 336 F.3d 811, 818–22 (9th Cir. 2003).

³⁶ *See A&M Recs, Inc. v. Napster, Inc.*, 239 F.3d 1004, 1014–19 (9th Cir. 2001); *Fox News Network, LLC v. Tveyes, Inc.*, 883 F.3d 169, 181 (2d Cir. 2018).

³⁷ *Stewart v. Abend*, 495 U.S. 207, 236 (1990); *see also* H.R. Rep. 94–1475 at 65–66 (1976) (fair use should be “adapt[ed]” to account for “rapid technological change”).

³⁸ 17 U.S.C. § 107.



of the Copyright Act, which includes four non-exclusive factors for courts to consider in determining whether a use is “fair”:

- (1) the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes;
- (2) the nature of the copyrighted work;
- (3) the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and
- (4) the effect of the use upon the potential market for or value of the copyrighted work.³⁹

As detailed below, these statutory factors, along with the many decades of case law interpreting them in the context of new technologies, confirm that it is “not an infringement” of copyright for a Generative AI model to “learn” by deriving statistical information from copyrighted texts, images, or other media. We address these factors in turn below.⁴⁰

1. Extracting Information To Create New Content Is Quintessentially Transformative⁴¹

One of the hallmarks of a transformative fair use in the context of technology is the use of a copyrighted work to create a useful, non-infringing output. This is the reason we have internet search engines,⁴² online book repositories,⁴³ and video game emulators.⁴⁴ Each of these technologies involves copying of one or many copyrighted works. The reason each is legitimate is that use is in service of a non-exploitive purpose: to extract information from the works and put that information to use by enabling the creation of new, non-infringing works, thereby “expand[ing] [the works’] utility.”⁴⁵

Extracting information from training data is the sole purpose and function of Generative AI model training. *See supra* Section II.B. As demonstrated above, the point of feeding data into an AI model is to enable that model to identify *patterns* across a broad number of texts, images, or other media. It is these patterns—not the training data’s protectable expression—on which the model relies. It is well-established that copyright law does not protect facts or the syntactical, structural, and linguistic information that training models are designed to extract. As the Second Circuit held in the seminal *Google Books* case, extracting information “about” a copyrighted work—like “statistical information” about “word frequencies, syntactic patterns, and thematic

³⁹ *Id.*

⁴⁰ Because Generative AI models can train on all manner of works, we do not discuss the second fair use factor, which examines “the nature of the copyrighted work.” 17 U.S.C. § 107(2).

⁴¹ This section addresses questions 8 and 8.1.

⁴² *See Kelly*, 336 F.3d at 818–22.

⁴³ *Google Books*, 804 F.3d at 217–18.

⁴⁴ *Sony Computer Entm’t, Inc. v. Connectix Corp.*, 203 F.3d 596, 603–08 (9th Cir. 2000).

⁴⁵ *Google Books*, 804 F.3d at 214.



markers”—in order to use that information for a further purpose is “the sort of transformative purpose” that “strongly favor[s] satisfaction of the first [fair use] factor.”⁴⁶

Moreover, the purpose of extracting this information is *not* to reproduce any aspect of the protected expression found in the training data, but rather to teach the model how language works, and then by extension allow the model to respond to user prompts with *wholly new* content. In this sense, the process behind Generative AI is similar to human learning. Just as a child learns language (words, grammar, syntax, sentence structure) by hearing everyday speech, bedtime stories, songs on the radio, and so on, a model “learns” language by being exposed—through training—to massive amounts of text from various sources. We learn by hearing others speak and reading the way others write—how words are put together cohesively, how word or punctuation placement can alter the meaning of a sentence, or even the different meanings of homonyms. Indeed, every author first benefited from learning language through the world around them, and then reading hundreds or thousands of existing books, consciously or subconsciously identifying patterns and concepts from those books, and then using those patterns and concepts to create their own works that compete in the general market for books. The fact that Generative AI accomplishes the same process through technology does not transform it from legitimate learning into a proscribed act of copyright infringement. Many courts that have considered that question have recognized that the creation of copies of copyrighted works (especially copies that are not perceivable to the public) in the course of technological development of non-infringing, competing products is protected by fair use.⁴⁷

The Supreme Court’s recent decision in *Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith*, further reinforces these fundamental fair use principles.⁴⁸ The Court in *Warhol* explained that fair use’s first factor proceeds by asking two questions: (1) first, whether the defendant’s use “has a further purpose or different character,” and (2) if not, whether there is a “justification for [the] copying.”⁴⁹ Significantly, the Court emphasized that each “use” must be assessed separately.⁵⁰ For that reason, it focused on the Warhol Foundation’s licensing of Warhol’s “Orange Prince” (which was based on a photograph of Prince by the plaintiff Lynn Goldsmith) to Conde Nast in 1984 to illustrate a story about the artist Prince, rather than Andy Warhol’s original creation of the painting.⁵¹ Viewed through that lens, the Court concluded that the Warhol Foundation’s use “share[d] substantially the same purpose” as the original, because Goldsmith’s purpose for creating the original photograph was to illustrate a magazine article about Prince in *Newsweek*, and that she had licensed that photograph to other magazines as well.⁵²

⁴⁶ *Google Books*, 804 F.3d at 209, 217 (expressing “no doubt” about this conclusion).

⁴⁷ See *Assessment Technologies of WI, LLC v. WIREdata, Inc.*, 350 F.3d 640, 644–45 (7th Cir. 2003); *Connectix*, 203 F.3d at 603–08; *Bateman v. Mnemonics, Inc.*, 79 F.3d 1532, 1539 n.18 (11th Cir. 1996); *Atari Games Corp. v. Nintendo of Am. Inc.*, 975 F.2d 832, 836–37, 843–44 (Fed. Cir. 1992); *Sega Enterprises Ltd. v. Accolade, Inc.*, 977 F.2d 1510, 1522–23 (9th Cir. 1992).

⁴⁸ 143 S. Ct. 1258 (2023).

⁴⁹ *Id.* at 1277.

⁵⁰ See *id.*

⁵¹ See *id.* at 1267.

⁵² *Id.* at 1273.



Because of that overlap in purpose, the “commercial nature” of the Warhol Foundation’s use “loom[ed] larger.”⁵³

As a result, the Court explained that the Warhol Foundation needed to present “some other justification for copying.”⁵⁴ As the Court explained, the scope of justifications that would qualify at this step is broad, ranging from parody or criticism,⁵⁵ to technological justifications like allowing “different programs to speak to each other” or making sure that “programmers are [] able to use their acquired skills.”⁵⁶ But because (as the Court found) no such justification applied to the Warhol Foundation’s use, the first factor weighed against fair use. *Id.* at 1287. Ultimately, however, Warhol reaffirmed that courts, in examining the “purpose” and the “justification” for the copying, should assess whether the use “furthers the goal of copyright, namely, to promote the progress of science and the arts, without diminishing the incentive to create.” *Warhol*, 143 S.Ct. at 1276.

Applying the *Warhol* rubric here only reinforces the degree to which training of AI models is transformative. The works that AI models are trained on—whether books, blog posts, photographs, videos, etc.—were created for expressive purposes. The purposes of AI developers’ use of those works is wholly distinct; models use training data not to copy their content or challenge authors’ ability to sell copies of their works, but rather to develop an entirely new and innovative service that, in turn, produces valuable new content—thereby vastly expanding the capacity for human creative productivity and the “progress of science and the useful arts.” U.S. CONST. Art. I, § 8, cl. 8. Given that wholly distinct purpose, the commercial nature of the use becomes less important. Moreover, there is an overwhelming justification for the copying; the ability to accurately distill the desired facts (whether about language or images or sounds) requires the ingestion of massive amounts of content that cannot reasonably be individually licensed. In the context of Generative AI, that justification is clear—to reject that justification would be to reject the thrust of the case law finding fair use on similar grounds.⁵⁷

The *Warhol* decision did not depart from that case law, but reinforced its continued vitality. For instance, it approvingly cited the *Google Books* case three times, thereby endorsing the essential holding of that case: that copying of “original copyrighted books [] to make available significant information about those books” constitutes a “transformative purpose.”⁵⁸ And just two years prior, the Court approved other cases holding that the creation of a new product was sufficient justification for “direct use of [] copyrighted material”—even if the defendant’s product competes with the original.⁵⁹

⁵³ *Id.* at 1285.

⁵⁴ *Id.* at 1280.

⁵⁵ See *id.* at 1276 (citing *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 588 (1994)).

⁵⁶ *Id.* at 1277 n.8 (citing *Google LLC v. Oracle Am., Inc.*, 141 S. Ct. 1183, 1203 (2021)).

⁵⁷ See, e.g., *Google Books*, 804 F.3d at 209, 217; *Connectix*, 203 F.3d at 603–08; *Accolade*, 977 F.2d at 1522–23; *iParadigms*, 562 F.3d at 638–40, 645.

⁵⁸ *Google Books*, 804 F.3d at 217.

⁵⁹ *Accolade*, 977 F.2d at 1522–23; see also *Oracle*, 141 S. Ct. at 1198–99 (citing *Accolade* with approval).



2. The Amount of the Information Extracted From Training Data Is Miniscule⁶⁰

The third fair use factor concerns “the amount and substantiality of the portion [of the work] used in relation to the copyrighted work as a whole.”⁶¹ “[A] finding of fair use is more likely when small amounts . . . are copied than when the copying is extensive,” because “[t]he larger the amount . . . of the original that is copied, the greater the likelihood that the secondary work might . . . diminish the original rights holder’s sales and profits.”⁶² As the Supreme Court has explained, this analysis focuses, in particular, on whether “the amount and substantiality of the portion used . . . [is] reasonable in relation to the purpose of the copying.”⁶³

The process of Generative AI training is meant to extract, relatively speaking, a miniscule amount of information from each piece of training data. To illustrate, recall the example discussed above, regarding the influence of a single sentence—*Susan’s aunt planted the flowers in the garden*—on a simple “word embeddings” process. As explained, that sentence, in isolation, has negligible value to the model as a whole. Rather, it is the *combination* of that sentence with billions of others that allows the model to identify the patterns of semantic meaning that it encodes in word vectors.⁶⁴ The “amount and substantiality” of the use of each individual piece of training data is negligible because each piece has a negligible influence on the model’s store of statistical information.

The analysis of the third fair use factor is thus more straightforward in this context than it was in other cases, like *Google Books*. There, the Court analyzed this factor by analyzing the “quantity of the copyrighted text” that a Google Books user could see in order to evaluate “the likelihood that those revelations could serve [] as an effective, free substitute for the purchase of the plaintiff’s book.”⁶⁵ Because the Google Books platform made it “at least difficult, and often impossible, for a searcher to gain access to more than a [small portion]” of the works at issue, the Court found that this factor weighed in favor of fair use.⁶⁶

Generative AI models, by contrast, are not designed to “re-present” *any* of the “protected aspects of the original work” to users.⁶⁷ Put another way, a user would not use Generative AI models to access and read a book that the models might have used for training purposes. The purpose of the models is to extract enough *statistical* information about language and abstract concepts to enable the creation of *new* content—not to capture and reproduce expressive material from the training data itself. Indeed, teams of machine-learning researchers using sophisticated “data extraction attack” methods have largely failed to extract expressive material from Generative

⁶⁰ This section addresses questions 8 and 8.4.

⁶¹ 17 U.S.C. § 107(3).

⁶² *Google Books*, 904 F.3d at 221.

⁶³ *Campbell*, 510 U.S. at 586.

⁶⁴ *See supra* Section II.B.

⁶⁵ 804 F.3d at 222.

⁶⁶ *Id.* at 222–23.

⁶⁷ *Id.* at 225.



AI models.⁶⁸ The fact that “a searcher cannot succeed, even after long extended effort . . . in revealing . . . what could usefully serve as a competing substitute for” any individual piece of training data suggests that the amount and substantiality of the information extracted from that data is minimal.⁶⁹

3. Generative AI Does Not Cause Market Harm⁷⁰

The fourth fair use factor examines “the effect of the use upon the potential market for or value of the copyrighted work.”⁷¹ Critically, however, the goal of the fourth factor is not to protect rightsholders from any form of “economic harm.”⁷² Simply put, not every harm to “the potential market for or value of a copyrighted work” is cognizable under the Copyright Act. For instance, the law authorizes the creation of unlicensed parodies, even though “a lethal parody, like a scathing theater review, kills demand for the original.”⁷³

Rightsholders point to several alleged harms, each of which is discussed below.

a. Harm To Hypothetical Licensing Markets

First, rightsholders have asserted that there is a potential market for use of their works at the input stage, as training data. They point to their willingness to accept so-called “licensing” fees for use of their works for AI training as evidence that unlicensed use causes a market harm under the fourth factor.⁷⁴ That position runs headlong into a “circularity” that courts and commentators have repeatedly warned of: “a potential market, no matter how unlikely, has always been supplanted in every fair use case, to the extent that defendant, by definition, has made some actual use of plaintiff’s work, which use could, in turn, be defined in terms of the relevant potential market.”⁷⁵

That danger of circularity is particularly acute here, where the claimed licensing market is not for licenses for the enjoyment of the original works’ expressive content (*e.g.*, a license to

⁶⁸ Nicholas Carlini et al., “Extracting Training Data from Diffusion Models” at 4–5 (Jan. 30, 2023), available at: <https://arxiv.org/pdf/2301.13188.pdf> (out of 175 million images generated in a sophisticated “two-stage data extraction attack,” only 109 were duplicates or near-duplicates of the training data); *see also* Nicholas Carlini, et al. “Quantifying Memorization Across Neural Language Models” at 3–4, 9 (Mar. 6, 2023), available at: <https://arxiv.org/pdf/2202.07646.pdf> (finding memorization rates of roughly 1% for models trained on de-duplicated datasets).

⁶⁹ *Google Books*, 804 F.3d at 222.

⁷⁰ This section addresses questions 8, 8.5, 10.1, 10.2, 11, 12, 13, and 14.

⁷¹ 17 U.S.C. § 107(4).

⁷² *Authors Guild, Inc. v. HathiTrust*, 755 F.3d 87, 99 (2d Cir. 2014) (internal quotation marks omitted).

⁷³ *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 591–92 (1994).

⁷⁴ *See* Cala Coffman, “Does the Use of Copyrighted Works to Train AI Qualify as a Fair Use?,” Copyright Alliance (Apr. 11, 2023), available at: <https://copyrightalliance.org/copyrighted-works-training-ai-fair-use/> (arguing that the fourth factor cuts against fair use because “many developers do not compensate copyright owners for the works used to train generative AI” through “AI training licenses”).

⁷⁵ 4 NIMMER ON COPYRIGHT § 13F.08; *see also Oracle*, 141 S. Ct. at 1207 (highlighting “the ‘danger of circularity posed’ by considering unrealized licensing opportunities because ‘it is a given in every fair use



redistribute e-books through an online shopping platform), or for the creation of derivative works (e.g., a license to create a motion picture based on a book). Instead, as explained, the relevant market is for use of the works to glean unprotectable facts and statistics from them, in service of a technology that is meant to create entirely new content. In analogous circumstances, courts have been clear: a copyright holder “cannot prevent others from entering fair use markets merely ‘by developing or licensing a market for . . . transformative uses of its own creative work.’”⁷⁶ Put simply, there is no legitimate market through which copyright holders can be compensated for the use of statistical facts extracted from their works, because those facts are “instantly available for public exploitation at the moment of publication.”⁷⁷

Indeed, it would be impossible for any market to develop that could enable AI developers to license all of the data their models need. *Ringgold v. Black Entertainment Television, Inc.*, 126 F.3d 70, 81 (2d Cir. 1997) (plaintiff must show harm to a “traditional, reasonable, or *likely to be developed* market for licensing her work” (internal quotation marks omitted) (emphasis added)). As discussed above, Generative AI models need not only a massive *quantity* of content, but also a large *diversity* of content. To be sure, it is possible that AI developers will strike deals with individual rightsholders, to develop broader partnerships or simply to buy peace from the threat of litigation. But those kinds of deals would provide AI developers with the rights to only a miniscule fraction of the data they need to train their models. And it would be impossible for AI developers to license the rights to other critical categories of works—like internet reviews and other examples of casual, vernacular text—both because it would be impossible to locate the owners of such works, and administratively impossible to negotiate licenses with each of them.

As the Office has previously noted, the possibility that the Copyright Act would make “productive and beneficial” uses of content impractical due to the impossibility of locating relevant rightsholders is “difficult to reconcile with the objectives of the copyright system.”⁷⁸ Moreover, even if such licensing were possible, it would make open source licensing of AI models all but impossible—no company could afford to pay licensing fees based on third-party uses of that company’s models, and even tracking how models were used would be impracticable. This would lead to the very results that Meta’s open source model is meant to prevent: consolidation of this revolutionary technology into the hands of a handful of incumbent technology companies, and opaque models that perform poorly and are more subject to bias.

case that plaintiff suffers a loss of some *potential* market if that potential is defined as the theoretical market for licensing the very use at bar.” (quoting 4 NIMMER ON COPYRIGHT § 13.05[A][4]); *Cambridge Univ. Press v. Patton*, 769 F.3d 1232, 1278 (11th Cir. 2014) (“[L]icensing poses a particular threat that the fair use analysis will become circular, and [rightsholders] may not head off a defense of fair use by complaining that every potential licensing opportunity represents a potential market for purposes of the fourth fair use factor.”); Pierre N. Leval, “Toward a Fair Use Standard,” 103 Harv. L. Rev. 1105, 1124 (1990) (“By definition every fair use involves some loss of royalty revenue because the secondary user has not paid royalties.”).

⁷⁶ *Bill Graham Archives v. Dorling Kindersley Ltd.*, 448 F.3d 605, 614-15 (2d Cir. 2006) (quoting *Castle Rock Entm’t Inc. v. Carol Publ’g Grp.*, 150 F.3d 132, 145 n.11 (2d Cir. 1998)).

⁷⁷ *Eldred*, 537 U.S. at 219.

⁷⁸ U.S. Copyright Office, Report on Orphan Works and Mass Digitization at 35–36 (2015).



b. Harm From Substantially Similar Outputs

Second, some claim market harm based on the fact that AI models, at the prompting of a user and despite the best efforts of AI developers, sometimes output content that is, in a copyright sense, substantially similar to something in the training set.

To be clear, Generative AI models are not designed to “reve[al] [] sufficiently significant portions of” the protected material in training data so “as to make available a significantly competing substitute.”⁷⁹ And while one might use a chatbot to recommend a book or look up a historical fact, none of these uses reproduce enough expressive material as to compete with the original works in a manner that implicate rightsholders’ interests.⁸⁰ The Second Circuit reached the same conclusion in *Google Books*, finding that the challenged platform’s ability to “convey [] historical fact[s]” to “satisfy the searcher’s [research] need[s]” did not implicate “interests that are [] protected by the copyright.”⁸¹ And while it is possible (at least in theory) for Generative AI to create works “of the same type” that compete in the overall market with the originals, this is not the kind of substitution that implicates the fourth fair use factor, which does not punish uses that “simply enable[] the copier to enter the market for works of the same type as the copied work.”⁸²

Under current doctrine, a model that predominantly creates non-infringing outputs still qualifies for fair use protection, whether or not the model *can* be manipulated to extract segments of protectable training data in rare instances. For instance, in the *Google Books* case, the Court noted that there were “surely instances in which a searcher’s need for access to a text will be satisfied by the [Google Books] snippet view [function], resulting in either the loss of a sale to that searcher, or reduction of demand on libraries for that title, which might have resulted in libraries purchasing additional copies.”⁸³ Nonetheless the court ruled that the fourth fair use factor favored Google, finding that “the possibility, or even the probability or certainty, of some loss of sales does not suffice to make the copy an effectively competing substitute.”⁸⁴

The same is true here—the possibility that a researcher might be able to use a sophisticated data extraction attack on a foundation model to uncover training data, does not make that model an “effectively competing substitute.” Nor is the mere fact that a handful of users might be able to extract large portions of the underlying work or otherwise harmful outputs from general purpose Generative AI tools a grounds to impose liability on those tools’ developers.⁸⁵

⁷⁹ *Google Books*, 804 F.3d at 223–24.

⁸⁰ See *HathiTrust*, 755 F.3d at 99 (noting that “[b]ook reviews” are “a paradigmatic fair use”); see also *Feist Publications Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 344–45 (1991) (facts not copyrightable).

⁸¹ 804 F.3d at 224.

⁸² *Accolade*, 977 F.2d at 1523.

⁸³ 804 F.3d at 224.

⁸⁴ *Id.*; see also *A.V. ex rel. Vanderhye v. iParadigms, LLC*, 562 F.3d 630, 639–40 (4th Cir. 2009) (explaining “[t]he question of whether a use is transformative does not rise or fall on whether the use perfectly achieves its intended purpose,” finding use fair even though system “is not fool-proof”).

⁸⁵ See *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 442 (1984) (creator of technology “capable of substantial noninfringing uses” not liable); see also *Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913, 957 (2005) (Breyer, J., concurring) (noting that this rule “provide[s]



c. *Harm From Competition In The Market For Expressive Works*

Third, many rightsholders more broadly focus on the possibility that the output of Generative AI models might in some sense “compete” with either the original works they were trained on or against the potential future output of authors, even when that output does not embody any expressive content of any particular work in the training set. But a use that enables creation of new, non-derivative works that might compete with the original does not result in cognizable market harm.⁸⁶

Furthermore, in evaluating market effects, a court must not only look to the potential harms to rightsholders, but also to “the public benefits the copying will likely produce.”⁸⁷ This includes asking whether the benefits are “related to copyright’s concern for the creative production of new expression,” and whether they are “comparatively important . . . when compared with dollar amounts likely lost.”⁸⁸ Here, those benefits are clear—Generative AI is a historically significant advance in scientific knowledge and human creativity, designed to produce new expression in an entirely novel way. We have barely started to scratch the surface of the benefits of Generative AI technology, and with more time the public benefits will become even clearer.

V. Additional Copyright-Based Restrictions Would Do More Harm Than Good⁸⁹

Given the impossibility of licensing massive amounts of data through bilateral voluntary negotiations, some rightsholders and industry groups have suggested that Congress pass legislation to implement a statutory or collective licensing framework. Others have proposed legislation that would require AI developers to disclose trade secret-protected information about the models and their training process, or impose other burdensome affirmative obligations.

The AI industry is in its infancy, and the United States has the enviable position of leading the world in AI innovation and development due to its long-standing and principled approach to copyright law, which has made this country both a creative and technological leader. Countries around the world have adopted express and broad text- and data-mining (TDM) or fair use exceptions, creating similarly enabling environments for technological advancement and investment.⁹⁰

entrepreneurs with needed assurance that they will be shielded from copyright liability as they bring valuable new technologies to market”).

⁸⁶ See, e.g., *Accolade*, 977 F.2d at 1523 (that a use “simply enables the copier to enter the market for works of the same type as the copied work” does not establish market harm); *Connectix*, 203 F.3d at 607 (because competing product was “a legitimate competitor in the market,” that plaintiff might suffer “some economic loss . . . as a result of this competition” does not establish market harm); see also *Oracle*, 141 S. Ct. at 1198–99, 1208 (approving *Accolade* and *Connectix*).

⁸⁷ *Google LLC v. Oracle Am., Inc.*, 141 S. Ct. 1183, 1206 (2021).

⁸⁸ *Id.*

⁸⁹ This section addresses questions 10.3, 10.4, and 11.

⁹⁰ See Copyright Act of Japan, Art. 30-4 (“A work may be exploited regardless of the method used . . . for the purpose of information analysis . . . [including] analysis of information pertaining to language, sound, images or other elements constituting said information extracted from a large number of works or other



Proposals for collective or statutory licensing, in particular, ignore the massive administrative problems that they would create. Indeed, the Register of Copyrights acknowledged these very problems in her recent testimony to Congress.⁹¹ The high-performing language models in existence today (like Llama) were trained on a large portion of the publicly available information published on the internet—*i.e.* billions of pieces of text from millions of individual websites, many of whom created that content without any concern about how that content would be used—the vast, vast majority of which lacks any rightsholder information. Imposing a first-of-its-kind licensing regime now, well after the fact, will cause chaos as developers seek to identify millions and millions of rightsholders, for very little benefit, given that any fair royalty due would be incredibly small in light of the insignificance of any one work among an AI training set.

Proposals that require developers to disclose trade-secret protected information—including, *e.g.*, the composition and organization of training datasets—are similarly flawed. They would accomplish little practical benefit, while simultaneously imposing burdensome reporting obligations that force AI developers to disclose sensitive and valuable information to the world, including to our foreign rivals.

VI. Copyright Law Does Not, And Should Not, Protect Styles⁹²

As the Office noted in its Notice of Inquiry, several artists and performers have expressed concerns “about generative AI systems’ ability to mimic their voices, likenesses, or styles.”⁹³ While these concerns should be taken seriously, they are not new. Performers have successfully used state right of publicity law for decades to protect against exploitive and misleading imitations of their voices.⁹⁴ The same laws have long provided a remedy for celebrities and other public

large quantities of information.”); Singapore Copyright Act 2021, §§ 243–44 (permitting creation of copies for “computational data analysis,” including “using a computer program to identify, extract and analyse information” or “using the work or recording as an example of a type of information or data to improve the functioning of a computer program in relation to that type of data”); Copyright Act of Korea Art. 35-3 (permitting use that “does not unreasonably prejudice an author’s legitimate interest without conflicting with the normal exploitation of works” and mirroring 17 U.S.C. § 107 factors); State of Israel, Minister of Justice, “Opinion: Uses of Copyright Materials for Machine Learning” at 3 (Dec. 18, 2022) (concluding that “the use of copyrighted materials for [machine learning] is permitted under existing copyright doctrines”); Copyright Act of Malaysia § 13(2)(a) (permitting acts done “by way of fair dealing including for purposes of research, private study, criticism, review or the reporting of news or current events” and accompanied by acknowledgements); Copyright Act of Taiwan Art. 65 (mirroring 17 U.S.C. § 107 factors).

⁹¹ See Oversight of the U.S. Copyright Office, Hearing Before the House Judiciary Subcommittee on Courts, Intellectual Property, and the Internet at 1:11:24–42 (Sept. 27, 2023) (statement by Shira Perlmutter, Register of Copyrights) (“[T]here are a lot of practical issues involved [in a collective licensing scheme] that need to be explored including how the license fees would be set, how it would be distributed, and, if it is, how it can be made feasible given the volume of works that would be involved.”).

⁹² This section addresses questions 30, 31, and 32.

⁹³ 88 Fed. Reg. 59945.

⁹⁴ See, *e.g.*, *Midler v. Ford Motor Co.*, 849 F.2d 460 (9th Cir. 1988); *Waits v. Frito-Lay, Inc.*, 978 F.2d 1093 (9th Cir. 1992).



figures whose likenesses are misappropriated for commercial purposes.⁹⁵ And to the extent the use causes confusion, federal unfair competition law provides yet further recourse.⁹⁶

Style appropriation, however, is another matter. Artistic and musical styles have never been protectable under federal or state law because the public use of such styles is a bedrock requirement for the artistic and expressive freedom guaranteed by the First Amendment.⁹⁷ As one court put it: “Intellectual (and artistic) progress is possible only if each author builds on the work of others . . . Every work uses scraps of thought from thousands of predecessors.”⁹⁸ The mere fact that AI technology may make it easier for individuals to create work that borrows others’ artistic styles is not a reason to depart from these basic principles. To the contrary, a new technology that enables more people to engage in the creative process and fulfill their creative potential is something to be celebrated, not restricted.

In any case, any attempt to impose liability on the unauthorized use of an artist’s “style” would make no sense as a practical matter. “If an . . . artist claimed broad protection for a style not associated with a particular work . . . it would be difficult, if not impossible, to determine the scope of protection.”⁹⁹ Courts would be forced to evaluate artists’ claims to ownership over certain styles against contrary evidence that they merely cobbled together pre-existing stylistic elements borrowed from their predecessors. Judges and juries would undertake the role of art critics, comparing a challenged work with the “amorphous style” asserted by the plaintiff to determine whether the former hewed too closely to the latter.¹⁰⁰ These questions would be even more diffuse and specious than those at issue in the recent series of musical-work infringement cases involving Marvin Gaye, Led Zeppelin, Taylor Swift, and Ed Sheeran. The next generation of artists and musicians—who will no doubt leverage AI in their creative pursuits—will find themselves facing expressive restrictions we have never before imposed. A chill would fall over our vibrant creative industries.

In short, none of these concerns justifies a drastic departure from current law. Meta understands the concerns of the artistic community. But AI is a new phenomenon in its very early days. There is no reason to believe that legal doctrines like trademark, unfair competition, and the right of publicity are inadequate to the task of regulating it. Over the past several decades, courts have carefully navigated the intersection between those laws, the Copyright Act, and the speech-protective principles enshrined in the First Amendment.¹⁰¹ The resulting legal framework reflects

⁹⁵ See, e.g., *Onassis v. Christian Dior-New York, Inc.*, 472 N.Y.S.2d 254 (N.Y. Super. Ct. 1984).

⁹⁶ See 15 U.S.C. § 1125(a).

⁹⁷ See, e.g., *Jewelry 10, Inc. v. Elegance Trading Co.*, No. 88-cv-1320, 1991 WL 144151, at *4 (S.D.N.Y. July 20, 1991) (Leval, J.) (“[S]tyle[s]” and “technique[s]” like “pointillism, fauve coloring, cubism,” or “psychedelic colors” are “ideas”).

⁹⁸ *Nash v. CBS, Inc.*, 899 F.2d 1537, 1540 (7th Cir. 1990) (Easterbrook, J.).

⁹⁹ 2 PATRY ON COPYRIGHT § 4:14.

¹⁰⁰ *Id.*

¹⁰¹ See, e.g., *Dastar Corp. v. Twentieth Century Fox Film Corp.*, 539 U.S. 23, 34–36 (2003) (discussing intersection between copyright law and the federal Lanham Act); *Zacchini v. Scripps-Howard Broad. Co.*, 433 U.S. 562, 574–75 (1977) (discussing the need to balance and weigh First Amendment interests against the right of publicity); *ProCD, Inc. v. Zeidenberg*, 86 F.3d 1447, 1453 (7th Cir. 1996) (Easterbrook, J.)



a delicate balance between a panoply of overlapping and, at times, contradictory interests, including prevention of marketplace confusion, protection of individual privacy and autonomy, and the societal benefits of free expression. The Office should not support new restrictions on style, voice, or likeness appropriation that would upend that balance.

* * * * *

Meta applauds the Office’s efforts to examine this new technology and understand its implications for the law and for society in general. In particular, we appreciate the time and attention it has given to understanding the perspective of technology companies like Meta. We look forward to continuing our work with you as you study this rapidly developing field.

(discussing Copyright Act preemption of state-law claims, which “prevent[s] states from giving protection” to subject matter that Congress “has decided should be in the public domain”).